

We claim:

1. In a network having a plurality of security perimeter routers, a method for determining packets to be discarded in response to a distributed denial-of-service (DDoS) attack, said method comprising the steps of:

5 confirming a DDoS attack at a network location using a plurality of packet attribute values aggregated from said routers;

computing an aggregate conditional probability measure for each packet entering said location based on selected attributes included within said packet from each router;

10 computing an aggregate cumulative distribution function (CDF) of scores based on said computed aggregate conditional probability measures;

determining a discarding threshold using said cumulative probability function; and

sending said discarding threshold to each of said routers.

15

2. The method of claim 1, wherein said step of computing an aggregate conditional probability measure further includes the steps of updating an individual marginal probability mass function and a joint probability mass function for attributes carried by each said packet.

20

3. The method of claim 1, further including the step of granting immunity to packets of a specified sub-type entering said location.

25

4. The method of claim 1, wherein said aggregate conditional probability measure is computed in accordance with the following equation:

$$CP(p) = \frac{\rho_n}{\rho_m} \cdot \frac{JP_n(A = a_p, B = b_p, C = c_p, \dots)}{JP_m(A = a_p, B = b_p, C = c_p, \dots)}$$

5. The method of claim 1, wherein said aggregate conditional probability measure is computed in accordance with the following equation:

$$CP(p) = \frac{\rho_n}{\rho_m} \cdot \frac{P_n(A=a_p)}{P_m(A=a_p)} \cdot \frac{P_n(B=b_p)}{P_m(B=b_p)} \cdot \frac{P_n(C=c_p)}{P_m(C=c_p)}$$

6. The method of claim 1, wherein said discarding threshold is calculated
5 using a load shedding algorithm, combined with an inverse lookup on the aggregate CDF of scores.

7. The method of claim 2, wherein said joint and marginal probability functions are maintained using iceberg-style histograms.

10 8. In a network comprising a centralized controller and a plurality of routers forming a security perimeter, a method for selectively discarding packets during a distributed denial-of-service (DDoS) attack over said network, comprising:

15 aggregating victim destination prefix lists and attack statistics associated with incoming packets received from said plurality of routers to confirm a DDoS attack victim;

aggregating packet attribute distribution frequencies for incoming victim related packets received from said plurality of security perimeter routers;

20 generating common scorebooks from said aggregated packet attribute distribution frequencies and nominal traffic profiles;

aggregating local cumulative distribution function (CDF) of the local scores derived from said plurality of security perimeter routers; and

25 providing, to each of said plurality of security perimeter routers, a common discarding threshold, said discarding threshold defining a condition in which an incoming packet may be discarded at said security perimeter.

9. The method of claim 8 wherein said aggregating local victim destination prefix lists and attack statistics of incoming packets comprises:

30 comparing measured attribute values for packet traffic sent to a particular destination to nominal traffic attribute values;

identifying increases in said measured attribute values over said nominal traffic attribute values.

10. The method of claim 9, wherein said confirming said victim of said DDoS
5 attack comprises determining if said identified increases for said measured
attribute values exceed respective predetermined thresholds.

11. The method of claim 8, wherein said local victim destination prefix list
and attack statistics comprise at least one of packets per second (pps), bits per
10 second (bps), flow counts, and flow rates of incoming packets.

12. The method of claim 8, wherein said aggregating packet attribute
distribution frequencies for incoming victim related packets comprises:

receiving packet attribute distribution frequencies from said plurality of
15 security perimeter routers, said packet attribute distribution frequencies
including incoming packet attribute information comprising at least one of IP
protocol-type values, packet size, source/destination port numbers,
source/destination IP prefixes, Time-to-Live (TTL) values, IP/TCP header
length, TCP flag combinations, use IP fragmentation, and incorrect packet
20 protocol checksums.

13. The method of claim 8, wherein said aggregating packet attribute
distribution frequencies for incoming victim related packets comprises:

receiving packet attribute distribution frequencies from said plurality of
25 security perimeter routers, said packet attribute distribution frequencies
including incoming packet attribute information comprising joint distribution of
the fraction of packets having various combinations of Time-to-Live (TTL) values
and source IP prefix, packet-size and protocol-type, and destination port
number and protocol-type.

30
14. The method of claim 13, wherein said receiving packet attribute
distribution frequencies comprises receiving iceberg-style histograms
comprising said incoming packet attribute information.

15. The method of claim 8, wherein said generating common scorebooks comprises:

computing partial scores of different attributes; and

5 computing a weighted sum of said partial scores to yield a logarithmic function of conditional legitimate probability for each incoming packet.

16. The method of claim 8 wherein said common discarding threshold comprises:

10 performing a load-shedding algorithm to determine a fraction (%_{PD}) of arriving suspicious packets required to be discarded; and

performing an inverse lookup on the aggregate CDF of scores.

15. The method of claim 16, where at each of said plurality of security perimeter routers, said method further comprises:

determining whether a score of an incoming packet is less than or equal to said discarding threshold;

discarding said incoming packet in an instance said score is less than or equal to said discarding threshold; and

20 forwarding said incoming packet for routing to destination in an instance said score is greater than to said discarding threshold.

25 18. A method for selectively discarding packets at a plurality of routers forming a security perimeter during a distributed denial-of-service (DDoS) attack over a network, each of said routers comprising the steps of:

sending victim destination prefix list and attack statistics associated with incoming packets to a centralized controller adapted to confirm a victim of said DDoS attack;

30 sending packet attribute distribution frequencies for incoming victim related packets;

receiving, from said centralized controller, common scorebooks formed by aggregated packet attribute distribution frequencies and nominal traffic profiles;

5 sending a local cumulative distribution function (CDF) of scores to said centralized controller; and

discarding incoming packets based on a commonly distributed discarding threshold defined by said centralized controller.

10 19. The method of claim 18, further including the step of classifying said incoming packets as being one of suspicious and non-suspicious packets based on a destination address of said incoming packet.

15 20. The method of claim 19 wherein said local victim destination prefix list and attack statistics comprise at least one of packets per second (pps), bits per second (bps), flow counts, and flow rates of incoming packets.

20 21. The method of claim 19 wherein said sending packet attribute distribution frequencies comprises monitoring packet attribute distribution frequencies including incoming packet attribute information comprising at least one of IP protocol-type values, packet size, source /destination port numbers, source/destination IP prefixes, Time-to-Live (TTL) values, IP/TCP header length, TCP flag combinations, use IP fragmentation, and incorrect packet protocol checksums.

25 22. The method of claim 21 wherein said packet attribute distribution frequencies are sent in a form of iceberg-style histograms.

23. The method of claim 20 wherein said sending a local cumulative distribution function (CDF) of scores comprises:

30 determining a predetermined number of incoming packets to monitor;

for each incoming packet of said predetermined number of incoming packets:

determining attribute scores from said received scorebooks; and
locally aggregating said scores; and

5 forming said CDF from said aggregated scores associated with said predetermined number of incoming packets.

24. The method of claim 19 wherein said commonly distributed discarding threshold comprises:

10 a fraction (%_{PD}) of arriving suspicious packets associated with an aggregated CDF from all of said routers.

25. The method of claim 23, wherein said discarding said incoming packets comprises:

15 determining whether a score of an incoming packet is less than or equal to said discarding threshold;
discarding said incoming packet in an instance said score is less than or equal to said discarding threshold; and
forwarding said incoming packet for routing to destination in an instance
20 said score is greater than to said discarding threshold.

26. In a network having a plurality of security perimeter routers and a centralized controller for determining packets to be dropped in regard to a potential distributed denial-of-service (DDoS) attack at a location within a packet

25 network, said centralized controller comprising:

means for aggregating a plurality of packet attribute values respectively received from said routers to confirm said attack at said location;

means for computing an aggregate conditional probability measure for each packet entering said location based on selected attributes included within said packet from each location;

30 means for computing an aggregate cumulative distribution function (CDF) based on said computed aggregate conditional probability measures;

means for determining a drop threshold based on access to said cumulative probability function;

means for sending said drop threshold to each of said routers, wherein said routers are adapted to pass through packets that exceed said determined drop threshold to said location.

27. In a network having a plurality of security perimeter routers and a centralized controller for determining packets to be dropped in regard to a potential distributed denial-of-service (DDoS) attack at a location within a packet network, said centralized controller comprising:

means for aggregating, local victim destination prefix lists and attack statistics associated with incoming packets received from a plurality of routers forming a security perimeter in said network to confirm a victim of said DDoS attack;

means for aggregating packet attribute distribution frequencies for incoming victim related packets received from said plurality of security perimeter routers;

means for generating common scorebooks from said aggregated packet attribute distribution frequencies and nominal traffic profiles;

means for aggregating local cumulative distribution function (CDF) of the local scores derived from said plurality of security perimeter routers; and

means for providing, to each of said plurality of security perimeter routers, a common discarding threshold, said discarding threshold defining a condition in which an incoming packet may be discarded at said security perimeter.

28. In a network having a plurality of security perimeter routers and a centralized controller for determining packets to be dropped in regard to a potential distributed denial-of-service (DDoS) attack at a location within a packet network, each of said routers comprising:

means for sending victim destination prefix lists and attack statistics associated with incoming packets to a centralized controller adapted to confirm a victim of said DDoS attack;

means for sending packet attribute distribution frequencies for incoming victim related packets;

means for receiving, from said centralized controller, common scorebooks formed by aggregated packet attribute distribution frequencies and nominal traffic profiles;

means for sending a local cumulative distribution function (CDF) of scores to said centralized controller; and

means for discarding incoming packets based on a commonly distributed discarding threshold defined by said centralized controller.